IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s):	R. H. Utley
Case:	3
Serial No.:	10/675,716
Filing Date:	September 30, 2003
Group:	2181
Examiner:	Richard B. Franklin
Title:	Processor with Input Data Block Discard Mechanism for Use in an Oversubscription Condition

---

AMENDED APPEAL BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicant (hereinafter "Appellant") hereby appeals the final rejection dated August 9, 2006 of claims 1-5 and 7-20 of the above-identified application.

REAL PARTY IN INTEREST

The present application is assigned of record to Agere Systems Inc. On April 2, 2007, the assignee Agere Systems Inc. completed a merger with LSI Logic Corporation, with the resulting entity being named LSI Corporation. LSI Corporation is the real party in interest.

RELATED APPEALS AND INTERFERENCES

There are no known related appeals or interferences.

STATUS OF CLAIMS

The present application was filed on September 30, 2003 with claims 1-20. Claim 6 was cancelled in an amendment filed on May 30, 2006. Claims 1-5 and 7-20 are currently pending in the application. Claims 1, 19 and 20 are the independent claims.

Each of claims 1-5 and 7-20 stands rejected under 35 U.S.C. §102(b). Claims 1-5 and 7-20 are appealed.

STATUS OF AMENDMENTS

There have been no amendments filed subsequent to the final rejection.

SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a processor that comprises a plurality of input ports, memory circuitry for storing data blocks associated with protocol data units and received by the processor at the input ports, and controller circuitry coupled to the memory circuitry and operative to discard certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor. A discarded data block indicator is generated for a given one of the input ports if a data block received at the given input port for a particular protocol data unit is discarded, and one or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator. Moreover, the controller circuitry is operative to maintain separate discarded data block indicators for respective ones of the plurality of input ports. In an illustrative embodiment described in the specification at, for example, page 6, line 23, to page 7, line 2, and page 8, lines 16-25, and shown in FIG. 2, processor 102 comprises a plurality of input ports 200, memory circuitry (input buffer 202 and PDU buffer 206) for storing data blocks associated with protocol data units and received by the processor at the input ports, and controller 204 coupled to the memory circuitry and operative to discard certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor.

Dependent claim 8 specifies that after a final data block of the particular protocol data unit is received at the given input port while the discarded data block indicator for the given input port is set to a first value (i.e., after a first data block of the particular protocol data unit has been discarded; see claim 2 upon which claim 8 depends), the controller circuitry is operative to enqueue the particular protocol data unit in a data unit buffer of the memory circuitry. In an illustrative embodiment described in the specification at page 8, lines 9-11 with reference to FIG. 2, after a final data block of the particular protocol data unit (PDU) is received at a given one of input ports 200-1...200-N while the discarded data block indicator for the given input port is set to the first value, the controller 204 enqueues the particular PDU in the PDU buffer 206.

Dependent claim 9 recites "the processor of claim 8 wherein the particular protocol unit is enqueued with an associated error flag set." Illustrative embodiments are described in the specification at page 8, lines 11-13, with reference to FIG. 2, and at page 10, lines 16-18, with reference to step 310 of method 300 shown in FIG. 3.

Dependent claim 10 recites "the processor of claim 9 being further operative to initiate a clean-up operation for the protocol data unit based at least in part on the associated error flag." Illustrative embodiments are described in the specification at page 8, lines 14-15, with reference to FIG. 2, and at page 10, lines 16-18, with reference to step 310 of method 300 shown in FIG. 3.

Dependent claim 13 recites "the processor of claim 1 wherein at least one of the input ports comprises a logical input port of the processor." In an illustrative embodiment described in the specification at, for example, page 6, lines 2-3, at least one of the input ports 200, as shown in FIG. 2, may comprise a logical input port of the processor.

Dependent claim 14 recites "the processor of claim 1 wherein at least one of the input ports comprises a physical input port of the processor." In an illustrative embodiment described in the specification at, for example, page 6, lines 1-2, at least one of the input ports 200, as shown in FIG. 2, may comprise a physical input port of the processor.

Independent claim 19 is directed to a method for use in a processor comprising a plurality of input ports for receiving data blocks associated with protocol data units. The method comprises a step of discarding certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the

processor. The method further comprises a step of generating a discarded data block indicator for a given one of the input ports if a data block received at a given input port for a particular protocol data unit is discarded. One or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator. Furthermore, separate discarded data block indicators are maintained for respective ones of the plurality of input ports. In an illustrative embodiment, described in the specification at, for example, page 10, lines 5-15 and shown in FIG. 3, a generating step 306 comprises setting a state bit for a given input port, and a discarding step 308 comprises discarding remaining blocks arriving for the input port for which the separate state bit has been set.

Independent claim 20 is directed to an article of manufacturing comprising machine-readable storage medium having program code stored thereon for use in a processor comprising a plurality of input ports for receiving data blocks associated with protocol data units. The program code, when executed in the processor, implements a step of discarding certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor. It further implements a step of generating a discarded data block indicator for a given one of the input ports if a data block received at the given input port for a particular protocol data unit is discarded. One or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator. Separate discarded data block indicators are maintained for respective ones of the plurality of input ports. In an illustrative embodiment described in the specification at, for example, page 9, lines 21-25, the program code, when executed on network processor 102 as shown in FIGS. 1, 2 and 4, implements generating step 306 and discarding step 308, as shown on FIG. 3. The machine-readable storage medium in this embodiment may be viewed as, for example, internal memory 104 or external memory 106, both associated with network processor 102, or a combination of such storage media.

The claimed invention provides a number of significant advantages over conventional arrangements. See the specification at, for example, page 3, line 26, to page 4, line 2

("Advantageously, the techniques of the invention in the illustrative embodiment respond to an oversubscription condition using an approach which ensures that the minimum number of PDUs are corrupted via discarded data blocks. As a result, processor performance and throughput are considerably improved.") and page 7, lines 9-13 ("This advantageously eliminates the need for random discarding of data blocks and thereby substantially reduces the number of PDUs for which data blocks are discarded. The invention in the illustrative embodiment thus ensures that the minimum number of PDUs are corrupted via discarded data blocks. As a result, processor performance and throughput are considerably improved.")

GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-5 and 7-20 stand rejected under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,764,641 (hereinafter "Lin").

ARGUMENT

Claims 1-5 and 7-20 stand rejected under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,764,641 (hereinafter "Lin").

Appellant asserts that Lin fails to teach or suggest all of the limitations in said claims for at least the reasons presented below.

Appellant initially notes that MPEP § 2131 specifies that a given claim is anticipated "only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference," citing Verdegaal Bros. v. Union Oil Co. of California, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Moreover, MPEP § 2131 indicates that the cited reference must show the "identical invention . . . in as complete detail as is contained in the . . . claim," citing Richardson v. Suzuki Motor Co., 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). For the reasons identified below, Appellant submits that the Examiner has failed to establish anticipation of claims 1-5 and 7-20 by Lin.

Claims 1-5, 7, 11, 12 and 15-20

Independent claim 1 is directed to a processor that comprises a plurality of input ports, memory circuitry for storing data blocks associated with protocol data units and received by the processor at the input ports, and controller circuitry coupled to the memory circuitry and operative to discard certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor. A discarded data block indicator is generated for a given one of the input ports if a data block received at the given input port for a particular protocol data unit is discarded, and one or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator. Moreover, the controller circuitry is operative to maintain separate discarded data block indicators for respective ones of the plurality of input ports.

Thus, in the invention as set forth in claim 1, each of the plurality of input ports has its own separate discarded data block indicator. In an illustrative embodiment of the invention, described in the specification at page 10, lines 5-11, the controller 204 in network processor 102 as shown in FIG. 2 maintains separate state bits for each of the N input ports 200-1 through 200-N in the set of input ports 200. See steps 304 through 308 of the flow diagram in FIG. 3. As indicated at page 7, lines 11-13, such an approach advantageously reduces the number of PDUs that are corrupted via discarded data blocks in an oversubscription condition, thereby improving processor throughput and performance. It is to be appreciated, of course, that these particular aspects of the illustrative embodiment are presented by way of example only, and are not to be construed as limitations of the claimed invention.

The Examiner in formulating the §102(b) rejection of claim 1 argues that each and every one of the above-noted limitations is met by Lin. Appellant respectfully disagrees. For example, in characterizing the Lin reference as allegedly meeting the limitation of independent claim 1 regarding controller circuitry operative to maintain separate discarded data block indicators for respective ones of the plurality of input ports, the Examiner relies primarily on column 6, lines 55-61, and FIG. 3B, steps 111-112, of Lin. See the final Office Action dated August 9, 2006, at page 4, lines 1-3. However, Appellant respectfully submits that the relied-upon portions of Lin

fail to anticipate the limitation at issue. The Lin reference, in column 6, lines 55-61, states the following, with emphasis supplied:

> If the queue length exceeds the EPD buffer threshold, the controller 18 discards the cell, since the controller cannot ensure that it will have available sufficient buffer space to hold the remaining cells of the packet as they arrive. The controller thus pre-empts what is expected to be a fragmented packet. The controller then sets <u>an EPD flag that is associated with the virtual circuit identified in the cell</u> (steps 111-112).

The Examiner apparently argues that the early packet discard (EPD) flag of Lin is anticipatory of the recited discarded data block indicator. However, claim 1 clearly indicates that <u>separate such discarded data block indicators are maintained for respective input ports of the processor</u>. The relied-upon portions of Lin indicate that EPD flags are set for particular <u>virtual circuits</u> that are identified in a corresponding cell. Thus, the EPD flags are not maintained for respective input ports of any processor.

It is important to note that the Examiner argues that the recited processor of claim 1 is met by switch 10 shown in FIG. 1 of Lin, and that the recited input ports are met by the input ports 12 of the switch 10 in FIG. 1 of Lin. See the final Office Action, at page 3, section 3, second paragraph. However, <u>Lin clearly does not teach to maintain separate discarded data block indicators for respective ones of the input ports 12 of switch 10</u>. As indicated previously, the EPD flags of Lin are maintained for virtual circuits, and such a disclosure clearly teaches away from maintenance of discarded data block indicators for individual ones of the input ports 12 of switch 10.

It should also be noted that Lin itself distinguishes the input ports 12 of switch 10 from virtual circuits that may traverse those ports. For example, Lin at column 4, lines 50-52, describes switch 10 as comprising a plurality of input ports 12 "that receive information over a network," and at column 6, lines 18-20, indicates that there may be multiple "switches along a virtual circuit." Thus, Lin describes a single virtual circuit as comprising multiple switches, each of which will comprise multiple input ports. In Lin, multiple input ports may therefore be associated with a single virtual circuit, and a given single input port can apparently be associated

with multiple virtual circuits. Such teachings are directly contrary to explicit limitations of independent claim 1.

The Evidence Appendix includes a number of references that illustrate that a given input port may be shared by multiple virtual circuits. See, e.g., Annabel Z. Dodd, The Essential Guide to Telecommunications 270 (3d ed. 2002) ("Multiple permanent virtual circuits can use one port."); Nathan J. Muller, Desktop Encyclopedia of Telecommunications 393 (3d ed. 2002) ("Oversubscription: Multiple permanent virtual circuits can share one access link, even . . . . oversubscribing the port."); J.A. Pecar et al., The New McGraw-Hill Telecom Factbook 599 (2d ed. 2000) ("[M]any virtual circuits can be carried through [one] port, each with its own destination.").

Accordingly, Appellant submits that the EPD flags set for particular virtual circuits as disclosed in Lin fail to meet the claim limitations relating to maintenance of separate discarded data block indicators for respective input ports of a processor.

The Examiner further contends, in an Advisory Action dated November 17, 2006, that "[t]he input port of the claimed invention is anticipated by the virtual circuit of Lin. The claims do not describe the structure of the port. Because of the lack of description of the input port, an interface point in a system meets the limitation of the claims. The virtual circuit of Lin is an interface point."

Appellant respectfully submits that it is well-settled law that "[d]uring examination . . . claim language should be read in light of the specification as it would be interpreted by one of ordinary skill in the art." See, e.g., In re Am. Acad. of Sci. Tech. Ctr., 367 F.3d 1359, 1364, 70 USPQ2d 1827, 1830 (Fed. Cir. 2004). Appellant further submits that the specification, at page 5, line 26, to page 6, line 6, describes the input port as something other than a mere "interface point in a system," with emphasis supplied:

> [T]he network processor 102 includes a set 200 of N input ports denoted 200-1, . . . 200-N. The input ports 200 may represent physical input ports of the processor, logical input ports of the processor, or any other type of input ports suitable for receiving data blocks associated with PDUs, as well as combinations of such input ports. For example, at least one of the input ports 200 may comprise a logical input port of the processor. The particular number N of input ports is

> purely arbitrary, and may take on any desired value supported by a given processor implementation.

Appellant notes that the Lin reference states as follows in its specification, at column 1, lines 28-35, with emphasis supplied:

> [T]he system establishes, <u>between the source station and the destination station, a particular route or "virtual circuit" over which the cells are to travel</u>. The source station transmits the cells, in order, over the virtual circuit and the cells arrive at the destination station in the same order. Cells from other packets that are traveling over a different, intersecting virtual circuit are typically interleaved with the cells, and thus, alter their time but not their order of arrival at the destination station.

and at lines 19-22, with emphasis supplied:

> <u>A destination station, after receiving all the cells associated with a packet in a plurality of transmissions</u>, reassembles the cells into the packet and provides the packet to the end system or application.

Appellant respectfully submits that the virtual circuit of Lin, which is described as a route over which cells travel between a source station and a destination station, is not equivalent to the input port of the present invention, which is described as suitable for receiving data blocks associated with PDUs.

Appellant further submits that dependent claims 2-5, 7, 11, 12 and 15-18 are believed allowable for at least the reasons identified above with regard to independent claim 1.

Independent claims 19 and 20 are believed allowable for reasons similar to those provided above in the context of claim 1.

<u>Claim 8</u>

Dependent claim 8 specifies that after a final data block of the particular protocol data unit is received at the given input port while the discarded data block indicator for the given input port is set to a first value (i.e., after a first data block of the particular protocol data unit has been discarded; see claim 2 upon which claim 8 depends), the controller circuitry is operative to

enqueue the particular protocol data unit in a data unit buffer of the memory circuitry. The Examiner contends that Lin discloses this limitation at column 6, line 62, to column 7, line 19. See the final Office Action at page 5, second paragraph. Appellant respectfully submits that the relied-upon portion of Lin not only fails to suggest this limitation but in fact teaches away by instead disclosing that received cells other than an EOP cell are discarded rather than enqueued. This is apparent from, for example, Lin at column 6, lines 62-64, where it states that the controller discards all of the received cells associated with the same packet until it receives the EOP cell.

Claim 9

Dependent claim 9 recites "the processor of claim 8 wherein the particular protocol unit is enqueued with an associated error flag set." Appellant submits that it is patentable for at least the reasons identified in reference to claims 1, 2 and 8 from which it depends. Additionally, Appellant submits that Lin fails to disclose the additional limitation related to an error flag. The Examiner apparently contends that this limitation is disclosed by the CLP bit of Lin, as disclosed in column 7, lines 9-19. See the final Office Action, at page 5, third paragraph. Appellant respectfully submits that the CLP bit of Lin is not an error flag as recited, but is in fact defined as indicating the loss-priority of the associated packet. See Lin at column 6, lines 8-19:

> As appropriate, the station also sets the CLP bit to 1 to indicate that the cells have low loss-priority, based on the application for which the cell is intended. The CLP bit set to 1 indicates that the cell may be readily discarded to avoid congestion, and a CLP bit that is instead set to 0 indicates that the cell should be retained, if possible. For example, the station sets the CLP bits to 1 in the cells that contain incidental information that is not specifically required by the end system, such as the fine tuning information for a video transmission. By default, the CLP bits are set to 0. The CLP bit 36 may be reset by switches along the virtual circuit, as appropriate, to correspond to a change in priority.

Claim 10

Dependent claim 10 recites "the processor of claim 9 being further operative to initiate a clean-up operation for the protocol data unit based at least in part on the associated error flag." Appellant submits that it is patentable for at least the reasons identified in reference to claims 1,

2, 8 and 9 from which it depends. Additionally, Appellant submits that Lin fails to disclose the additional limitation related to a clean-up operation. Indeed, the Examiner likewise fails to locate any support in Lin for this claim limitation, instead combining this rejection with that of claim 9 and basing both rejections on a portion of Lin dealing with the setting of a CLP bit upon receipt of an EOP cell. See the final Office Action, at page 5, third paragraph ("As per claim 9-10, Lin also teaches wherein the particular protocol data unit is enqueued with an associated error flag set (Figure 3A Item 103 "CLP=0"; Col 7 Lines 9-19).") There is no mention whatsoever of a clean-up operation in the relied-upon portion of Lin.

Claim 13

Dependent claim 13 recites "the processor of claim 1 wherein at least one of the input ports comprises a logical input port of the processor." Appellant respectfully submits that this additional claim limitation is not met for at least the reasons discussed above in reference to the limitation in claim 1 (i.e., that Lin fails to teach the input ports of claim 1). Appellant respectfully further submits that Lin fails to disclose, or even mention, logical input ports. The Examiner, in formulating this rejection, relies upon FIG. 1, item 12 of Lin and column 3, line 64, to column 4, line 8. See the final Office Action, at page 6, second paragraph. Appellant respectfully submits that FIG. 1, item 12 of Lin depicts a physical input port and that column 3, line 64, to column 4, line 8 of Lin not only does not refer to FIG. 1 at all (much less describing item 12 thereof), but is in fact completely inapposite to the limitation of claim 13:

> Each time a cell is discarded, whether for a UPC/NPC violation, because of an exceeded buffer limit or CLP threshold, the switch invokes its I-TPD scheme and discards all of the remaining cells of the same packet, except the EOP cell. It thus discards all the cells it receives over the virtual circuit identified in the first discarded cell until it receives the EOP cell. Assuming the queue has not reached its queue limit when the switch receives the EOP cell, the switch sets the priority of the EOP cell to high and passes it to the buffer. This increases the probability that the packet boundaries will be retained through the succeeding switches in the virtual circuit.

Claim 14

Dependent claim 14 recites "the processor of claim 1 wherein at least one of the input ports comprises a physical input port of the processor." Appellant submits that the virtual circuit of Lin is not equivalent to the recited physical input port for reasons similar to those discussed above in reference to claim 1 (i.e., that it is not equivalent to any processor input port). Additionally, Appellant submits that a virtual circuit cannot be equivalent to a physical input port, given that "virtual" by definition implies "not physical." See, e.g., In re Styleclick.com Inc., 58 USPQ2d 1523, 1526 (TTAB 2001):

> The dictionary evidence shows the term "virtual" defined as follows: "Not real. The term virtual is popular among computer scientists and is used in a wide variety of situations. In general, it distinguishes something that is merely conceptual from something that has physical reality." PC Webopaedia (1998). We take judicial notice of these other listings for the term: "Not physical. Exists in the software only or in the imagination of the machine." net.speak—the internet dictionary (1994); "Used generally to describe something without a physical presence or is not what it appears to be. Virtual reality, for example, is made up of computer-generated images and sounds rather than actual objects." The Computing Dictionary (1996); and "conceptual rather than actual, but possessing the essential characteristics of a real function." The Illustrated Dictionary of Microcomputers (3rd ed. 1990).

Furthermore, Appellant submits that the references cited above in regard to claim 1, which demonstrate that a virtual circuit is not considered to be an input port by those skilled in the art, further indicate that a virtual circuit is not considered to be a physical input port.

In view of the above, Appellant believes that claims 1-5 and 7-20 are in condition for allowance, and respectfully requests the withdrawal of the §102(b) rejection.

Respectfully submitted,

Date: April 19, 2007

Joseph B. Ryan
Attorney for Applicant(s)
Reg. No. 37,922
Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, NY 11560
(516) 759-7517

CLAIMS APPENDIX

1. A processor comprising:

a plurality of input ports;

memory circuitry for storing data blocks associated with protocol data units and received by the processor at the input ports; and

controller circuitry coupled to the memory circuitry and operative to discard certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor;

wherein a discarded data block indicator is generated for a given one of the input ports if a data block received at the given input port for a particular protocol data unit is discarded; and

wherein one or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator; and

wherein the controller circuitry is operative to maintain separate discarded data block indicators for respective ones of the plurality of input ports.

2. The processor of claim 1 wherein the controller circuitry sets the discarded data block indicator for the given input port to a first value when a first data block of the particular protocol data unit is discarded.

3. The processor of claim 2 wherein the controller circuitry is configured to automatically discard any remaining data blocks of the particular protocol data unit that are received at the given input port while the discarded data block indicator is set to the first value.

4. The processor of claim 1 wherein the discarded data block indicator for the given input port comprises a single bit.

5. The processor of claim 4 wherein the single bit being at a first logic level indicates that at least one data block received at the given input port has been discarded for a corresponding protocol data unit, and the single bit being at a second logic level indicates that no data block received at the given input port has yet been discarded for the corresponding protocol data unit.

6. (Canceled)

7. The processor of claim 1 wherein a given one of the discarded data block indicators indicates whether or not at least one data block received at the corresponding input port has been discarded.

8. The processor of claim 2 wherein after a final data block of the particular protocol data unit is received at the given input port while the discarded data block indicator for the given input port is set to the first value, the controller circuitry is operative to enqueue the particular protocol data unit in a protocol data unit buffer of the memory circuitry.

9. The processor of claim 8 wherein the particular protocol data unit is enqueued with an associated error flag set.

10. The processor of claim 9 being further operative to initiate a clean-up operation for the protocol data unit based at least in part on the associated error flag.

11. The processor of claim 1 wherein the oversubscription condition is overcome by discarding only data blocks received at the given input port, and associated with the particular protocol data unit.

12. The processor of claim 1 wherein the received protocol data units are associated with frame-based data.

13. The processor of claim 1 wherein at least one of the input ports comprises a physical input port of the processor.

14. The processor of claim 1 wherein at least one of the input ports comprises a logical input port of the processor.

15. The processor of claim 1 wherein the protocol data unit comprises a packet.

16. The processor of claim 1 wherein the processor is configured to provide an interface for communication of the received protocol data units between a network and a switch fabric.

17. The processor of claim 1 wherein the processor comprises a network processor.

18. The processor of claim 1 wherein the processor is configured as an integrated circuit.

19. A method for use in a processor comprising a plurality of input ports for receiving data blocks associated with protocol data units, the method comprising the steps of:

discarding certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor; and

generating a discarded data block indicator for a given one of the input ports if a data block received at the given input port for a particular protocol data unit is discarded;

wherein one or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator; and

wherein separate discarded data block indicators are maintained for respective ones of the plurality of input ports.

20. An article of manufacture comprising a machine-readable storage medium having program code stored thereon for use in a processor comprising a plurality of input ports for

receiving data blocks associated with protocol data units, the program code when executed in the processor implementing the steps of:

discarding certain ones of the data blocks received at the input ports in an oversubscription condition in which the received data blocks exceed a designated capacity of the processor; and

generating a discarded data block indicator for a given one of the input ports if a data block received at the given input port for a particular protocol data unit is discarded;

wherein one or more additional data blocks received at the given input port for the particular protocol data unit are discarded based at least in part on the discarded data block indicator; and

wherein separate discarded data block indicators are maintained for respective ones of the plurality of input ports.

EVIDENCE APPENDIX

The references contained herein were made of record in a Supplemental Information Disclosure Statement filed on January 2, 2007 and entered by the Examiner. Although the Supplemental Information Disclosure Statement did not comply with 37 CFR § 1.97(e), Appellant respectfully submit that these documents should nonetheless be considered under MPEP §§ 609.05(c) ("To the extent that a document is submitted as evidence directed to an issue of patentability raised in an Office action, and the evidence is timely presented, applicant need not satisfy the requirements of 37 CFR 1.97 and 37 CFR 1.98 in order to have the examiner consider the information contained in the document relied on by applicant.") and 1205.02 ("If in his or her brief, appellant relies on some reference, he or she is expected to provide the Board with a copy of it in the evidence appendix of the brief."), as well as 37 CFR 41.37(c)(1).

The three references contained in this appendix are:

1) ANNABEL Z. DODD, "The Essential Guide to Telecommunications," Third Edition, Prentice Hall PTR, page 270, Upper Saddle River, NJ.

2) NATHAN J. MULLER, "Desktop Encyclopedia of Telecommunications," Third Edition, McGraw-Hill, page 393.

3) J.A. PECAR ET AL., "The New McGraw-Hill Telecom Factbook," Second Edition, page 599.

The Essential Guide to Telecommunications

3rd Edition

Annabel Z. Dodd

Prentice Hall PTR, Upper Saddle River, NJ 07458
http://www.phptr.com

Frame Relay for Transmitting Voice

Organizations use Frame Relay to replace private lines for voice traffic between sites. Voice over Frame Relay is improving but is not as good as the voice transmitted on the standard public network.

The technologies used to transmit voice on Frame Relay are voice compression and silence suppression. Silence is suppressed so that pauses between words are used to transmit data and voice from other users. In addition, the voice itself is compressed, or made smaller, so that it does not require as much network capacity. Finally, voice traffic is given a high priority so that delay in voice conversations is minimized. The ATM switches (discussed later) located in core Frame Relay networks are able to prioritize packets containing voice traffic.

If the Frame Relay network becomes highly congested, voice quality can be degraded because packets are dropped or delayed. For this reason, some customer equipment has the ability to monitor traffic levels on the Frame Relay network and send it over the public switched network if there is congestion.

Frame Relay Pricing—Ports, Circuits and Committed Information Rate

Frame Relay service is priced at fixed monthly fees based on the following four elements, plus the cost of the telephone line used to connect each site to the Frame Relay service:

- **The permanent virtual circuit** (PVC) is a logical predefined path or link through a carrier's network. For example, if San Francisco and Tucson sites need to exchange data, the carrier defines a permanent virtual circuit between these two locations. PVCs are priced at fixed monthly fees.
- **The switched virtual circuit** (SVC): Unlike permanent virtual circuits, SVC charges are based on usage. Temporary connections are set up between points on a Frame Relay network. SVCs can be used to carry voice traffic if volumes are low. Thus, users only pay for what they use instead of incurring fixed monthly fees associated with permanent virtual circuits.
- **The Frame Relay port** is the entry point, on a Frame Relay provider's switch, to the Frame Relay network. Multiple permanent virtual circuits can use one port. Ports are available in variable speeds such as T-1, 56 Kbps, 256 Kbps and 512 Kbps.
- **The committed information rate** (CIR) is the minimum guaranteed number of bits-per-second throughput, typically half the capacity of the port

The *New* McGraw-Hill Telecom Factbook

Second Edition

Joseph A. Pecar
David A. Garbin

McGraw-Hill
New York San Francisco Washington, D.C. Auckland Bogotá
Caracas Lisbon London Madrid Mexico City Milan
Montreal New Delhi San Juan Singapore
Sydney Tokyo Toronto

guarantees of average throughput between specific end-points by establishing *virtual circuits* between them and by associating a *committed information rate* (CIR) with each virtual circuit (see Chapters 7 and 14).

Since frame relay combines attributes of a shared service (frame relay switches are shared among customers) with aspects of a dedicated service, its pricing structure reflects elements of both. Users access frame relay service through *ports*. A port terminates an access circuit and usually matches the access circuit's bit rate. As described in the example in Chapter 14, many virtual circuits can be carried through the port, each with its own destination and CIR. The sum of the CIRs may exceed the port bandwidth, since not all virtual circuits may be active at the same time. However, the total bit rate at any one time may not exceed the port capacity. It is not surprising, then, that a port charge is one of the pricing elements of frame relay service, and that the charge depends on the port bandwidth.

For connectionless data services, this port charge is the only charge. The additional functionality that frame relay provides on each permanent virtual circuit (the guarantee of average throughput equal to the purchased CIR) is paid for through a PVC charge that depends on the CIR ordered. In contrast with dedicated circuits, the prices for PVCs are flat rated and do not depend on the mileage between end-points or the location of the end-points. Since they are implemented in software, PVCs can be offered in finer steps and at bit rates lower than 64 kb/s, making them much easier to match to requirements than dedicated transport services. As noted previously, a PVC is also much more flexible in allowing the actual bit rate to fluctuate above or below the CIR.

Figure 15.8 shows a price schedule for frame relay service taken from the federal government's FTS2001 contract, a large multiservice contract providing a wide range of services to all federal agencies. The PVCs in this example are full-duplex virtual circuits (the CIR is guaranteed in both directions over the PVC). In some tariffs, PVCs are specified and priced as simplex PVCs (the CIR is guaranteed in only one direction). A separate PVC of a different CIR could be provisioned between the same two points in the opposite direction.

Figure 14.10 provides a direct comparison between a frame relay and a dedicated-circuit solution to a four-node customer configuration. The additional capabilities provided by the frame relay service are discussed there. The table below compares the total monthly costs of the two solutions using dedicated-circuit prices from FTS2001 and assuming 600 miles between the network POPs. In the dedicated-

Desktop Encyclopedia of Telecommunications

Nathan J. Muller

Third Edition

McGraw-Hill
New York Chicago San Francisco
Lisbon London Madrid Mexico City Milan
New Delhi San Juan Seoul Singapore
Sydney Toronto

as much as 50 percent, improving network throughput. Although the frame relay network can detect errors, it does not correct them. Bad frames are simply discarded. When the receiving device detects corrupt or missing frames, it can request a retransmission from the originating device, whereupon the appropriate frames are sent again.

Advantages of Frame Relay

The most compelling advantages of a carrier-provided frame relay service include:

- *Improved throughput/low delay* Frame relay service uses high-quality digital circuits end to end, making it possible to eliminate the multiple levels of error checking and error control. The result is higher throughput and less delay compared to legacy packet-switched networks like X.25.
- *Any-to-any connectivity* Any node connected to the frame relay service can communicate with any other node via predefined permanent virtual circuits (PVCs) or dynamically via switched virtual circuits (SVCs).
- *No long distance charges* Since frame relay is offered as a service over a shared network, the need for a highly meshed private line network is eliminated, for substantial cost savings. There are no distance-sensitive charges with frame relay, as there is with private lines.
- *Oversubscription* Multiple permanent virtual circuits can share one access link, even exceeding the port speed of the frame relay switch. In oversubscribing the port, multiple users can access the frame relay network—but not all at the same time—eliminating the cost of multiple private-line circuits and their associated customer premises equipment (CPE), for further cost savings.
- *Higher speeds* Whereas X.25 tops out at 56 Kbps, frame relay service supports transmission speeds up to 44.736 Mbps. If the frame relay switches in the network support Frame Relay Forum Implementation Agreement 14 (FRF 14), speeds at the OC-3 rate of 155 Mbps and the OC-12 rate of 622 Mbps over fiber backbones are possible.
- *Simplified network management* Customers have fewer circuits and less equipment to monitor. In addition, the carrier provides proactive monitoring and network maintenance 24 hours a day.
- *Intercarrier connectivity* Frame relay service is compatible between the networks of various carriers, through network-to-network interfaces (NNIs), enabling data to reach locations not served by the primary service provider.

<u>RELATED PROCEEDINGS APPENDIX</u>

None